

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s): Chen et al.
Docket No.: YO999-172
Serial No.: 09/345,238
Filing Date: June 30, 1999
Group: 2626
Examiner: Qi Han

Title: Method and Apparatus for Tracking Speakers in an Audio Stream

CORRECTED APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicants hereby submit this Corrected Appeal Brief in response to the Notification of Non-Compliant Appeal Brief dated April 9, 2008. The original Appeal Brief was submitted on March 18, 2008 to appeal the final rejection dated December 18, 2007, of claims 1 through 35 of the above-identified patent application.

REAL PARTY IN INTEREST

The present application is assigned to IBM Corporation, as evidenced by an assignment recorded on September 28, 1999 in the United States Patent and Trademark Office at Reel 010271, Frame 0023. The assignee, IBM Corporation, is the real party in interest.

RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

STATUS OF CLAIMS

Claims 1 through 35 are pending in the above-identified patent application. Claims 1-5, 8, 10-14, 16-19, 21-26 and 28-35 remain rejected under 35 U.S.C. §102(b) as being anticipated by Chen et al., "Speaker, Environment and Channel Change Detection and Cluster via the Bayesian Information Criterion," Proc. of the DARPA Broadcast News Workshop (Feb. 1998), hereinafter, referred to as "Chen." In addition, claims 6, 7, 9, 20, and 27 remain rejected under 35 U.S.C. §103(a) as being unpatentable over Chen in view of well known prior art and claim 15 remains rejected under 35 U.S.C. §103(a) as being unpatentable over Chen in view of Kleider et al. (United States Patent No. 5,930,748). Claims 1, 16, 23, and 30-35 are being appealed.

STATUS OF AMENDMENTS

There have been no amendments filed subsequent to the final rejection.

SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 requires a method for tracking a speaker in an audio source (page 5, lines 12-26; and page 8, lines 16-19), said method comprising the steps of: identifying potential segment boundaries in said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and clustering homogeneous segments from said audio source (page 8, lines 1-6) substantially concurrently with said identifying step (page 2, lines 16-17; FIG. 1: 200 and FIG. 2: 200).

Independent claim 16 requires a method for tracking a speaker in an audio source (page 5, lines 12-26; and page 8, lines 16-19), said method comprising the steps of: identifying potential segment boundaries in said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and clustering segments from said audio source (page 8, lines 1-6) corresponding to the same speaker substantially concurrently with said identifying step (page 2, lines 16-17, and page 5, lines 16-18; FIG. 1: 200 and FIG. 2: 200).

Independent claim 23 requires a method for tracking a speaker in an audio source (page 5, lines 12-26; and page 8, lines 16-19), said method comprising the steps

of: identifying potential segment boundaries during a pass through said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and clustering segments from said audio source (page 8, lines 1-6) corresponding to the same speaker during said same pass through said audio source (page 2, lines 16-17; FIG. 1: 200 and FIG. 2: 200).

5 Independent claim 30 requires a system (FIG. 1: 100) for tracking a speaker in an audio source (page 5, lines 12-26; and page 8, lines 16-19), comprising: a memory (FIG. 1: 120) that stores computer-readable code; and a processor (FIG. 1: 110) operatively coupled to said memory, said processor configured to implement said computer-readable code, said computer-readable code configured to: identify potential
10 segment boundaries in said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and cluster homogeneous segments from said audio source (page 8, lines 1-6) substantially concurrently with said identification of segment boundaries (page 2, lines 16-17; FIG. 1: 200 and FIG. 2: 200).

Independent claim 31 requires an article of manufacture, comprising: a
15 computer readable medium having computer readable code means embodied thereon, said computer readable program code means comprising: a step to identify potential segment boundaries in said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and a step to cluster homogeneous segments from said audio source (page 8, lines 1-6) substantially concurrently with said identification of segment boundaries (page 2,
20 lines 16-17; FIG. 1: 200 and FIG. 2: 200).

Independent claim 32 requires a system (FIG. 1: 100) for tracking a speaker in an audio source (page 5, lines 12-26; and page 8, lines 16-19), comprising: a memory (FIG. 1: 120) that stores computer-readable code; and a processor (FIG. 1: 110) operatively coupled to said memory, said processor configured to implement said
25 computer-readable code, said computer-readable code configured to: identify potential segment boundaries in said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and cluster segments from said audio source (page 8, lines 1-6) corresponding to the same speaker substantially concurrently with said identification of segment boundaries (page 2, lines 16-17, and page 5, lines 16-18; FIG. 1: 200 and FIG. 2: 200).

30 Independent claim 33 requires an article of manufacture, comprising: a computer readable medium having computer readable code means embodied thereon,

said computer readable program code means comprising: a step to identify potential segment boundaries in said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and a step to cluster segments from said audio source (page 8, lines 1-6) corresponding to the same speaker substantially concurrently with said identification of segment boundaries (page 2, lines 16-17, and page 5, lines 16-18; FIG. 1: 200 and FIG. 2: 200).

Independent claim 34 requires a system (FIG. 1: 100) for tracking a speaker in an audio source (page 5, lines 12-26; and page 8, lines 16-19), comprising: a memory (FIG. 1: 120) that stores computer-readable code; and a processor (FIG. 1: 110) operatively coupled to said memory, said processor configured to implement said computer-readable code, said computer-readable code configured to: identify potential segment boundaries during a pass through said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and cluster segments from said audio source (page 8, lines 1-6) corresponding to the same speaker during said same pass through said audio source (page 2, lines 16-17, and page 5, lines 16-18; FIG. 1: 200 and FIG. 2: 200).

Independent claim 35 requires an article of manufacture, comprising: a computer readable medium having computer readable code means embodied thereon, said computer readable program code means comprising: a step to identify potential segment boundaries during a pass through said audio source (page 7, lines 10-24; FIG. 1: 300 and FIG. 3: 300); and a step to cluster segments from said audio source (page 8, lines 1-6) corresponding to the same speaker during said same pass through said audio source (page 2, lines 16-17; FIG. 1: 200 and FIG. 2: 200).

STATEMENT OF GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1, 16, 23, and 30-35 are rejected under 35 U.S.C. § 102(b) as being anticipated by Chen et al.

ARGUMENT

Independent Claims 1, 16, 23, and 30-35

Independent claims 1, 16, 23 and 30-35 are rejected under 35 U.S.C. § 102(b) as being anticipated by Chen et al.

In the Office Action dated August 27, 2002, the Examiner asserted that Chen discloses speaker, environment and channel change detection and clustering via the Bayesian Information Criterion for segmenting the audio stream into homogeneous regions according to speaker identity, environmental condition and channel condition and clustering speech segments into homogeneous clusters according to speaker identity, environmental condition and channel (citing page 1, paragraph 2) which reads on the claimed "method of tracking a speaker in an audio source, said method comprising the steps of identifying potential segment boundaries in said audio source; and clustering homogeneous segments from said audio source substantially concurrently with said identifying step."

In the Response to Office Action dated December 26, 2002, Appellants submitted that, while Chen discloses segmenting an audio stream into homogeneous regions and clustering speech segments into homogeneous clusters, the audio stream is *first* segmented and *then* clustered. Appellants noted, as further evidence that the clustering in Chen is performed only after the audio stream has been segmented, that Section 4.1 indicates that each segment is compared to all other segments before clustering is finalized. In addition, Section 4.2, first paragraph indicates that the data set consists of an audio file that has been "hand-segmented into 824 short segments."

In the Office Action dated March 7, 2003, the Examiner notes that the prior art cites that "our segmentation algorithm can successfully detect acoustic changes" (Chen: abstract) and that "we first examine whether our detected change points were true." (Chen: Section 3.3, paragraph 3.) The Examiner asserts that this suggests that Chen not only employs its own segmenting mechanism, but is also capable of combining segmentation with clustering "substantially concurrently."

The Examiner also asserts that Chen suggests that the clustering does not need completely segmented data, such that a clustering process may be combined with a segmenting process together substantially concurrently, since Chen discloses that "it is also clear that our criterion can be applied to top-down methods." (Chen: Section 4.1, paragraph 4.)

The Examiner further asserts that a clustering step can be inserted in the segmentation loop, in Chen, Section 3.2, paragraph 1, and that Chen is capable of combining segmentation and clustering since the segmentation and clustering algorithms are based on the BIC algorithm and since equations (2), (3), and (8) have no limitation for combining segmentation and clustering.

Appellants acknowledge that Chen employs its own segmenting mechanism, but find no indication of or suggestion to perform segmentation and clustering “substantially concurrently” in the cited text. Appellants note that the Examiner asserts that Chen is *capable of* this, but does not assert that Chen suggests or discloses combining segmentation with clustering substantially concurrently.

Appellants also note that, in the top-down method, a hypothesis is made regarding the number of clusters. Then, a test is made to determine if the number of clusters hypothesized actually “fits” the data. Alternatively, in the bottom-up method, the number of clusters is determined from the data. Thus, the capability to utilize a top-down method does not suggest that segmentation is performed substantially concurrently with the clustering process.

Regarding the final assertion made by the Examiner, Appellants also note that, whether or not Chen is *capable of* combining segmentation and clustering, there is no disclosure or suggestion to do so.

Thus, Chen does not disclose or suggest a “method of tracking a speaker in an audio source, said method comprising the steps of identifying potential segment boundaries in said audio source; and clustering homogeneous segments from said audio source substantially concurrently with said identifying step,” as required by independent claims 1, 16, 30, 31, 32 and 33 of the present invention. Similarly, independent claims 23, 34 and 35 require that the segmentation and clustering are performed on the “same pass” through said audio source.

Response to Examiner’s Answer dated December 17, 2003

In the Examiner’s Answer dated December 17, 2003, the Examiner states that it is believed that the limitation “substantially concurrently” has no patentable weight, because the Applicant does not have any clear definition and/or description in the claim or in the specification about this limitation, and does not give any conditions to

apply this limitation. The Examiner also asserts that the prior art explicitly and/or implicitly discloses all the limitations regarding claim 1, including the limitation of “substantially concurrently,” based on the interpretation of the claim language and the understanding (of the) prior art teachings. In particular, the Examiner asserts that the performance of the two steps (segmentation and clustering) may be associated with many time related factors, including computing speed, sample rate, and total stream size.

The Examiner further asserts that the fact that the clustering in Chen is performed only after the audio stream has been segmented and that each segment is compared to all other segments before clustering is finalized is not relevant to claim 1 since claim 1 does not recite these limitations.

The Examiner also notes that one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references.

Regarding the Examiner’s assertion that the limitation “substantially concurrently” has no patentable weight, Appellants note that the word “substantially” has a well known and well understood definition in claim language. Its meaning is sufficiently clear in the teachings of the specification such that a person of ordinary skill in the art would understand the limitation without the need to apply conditions.

Regarding the Examiner’s assertion that the prior art explicitly and/or implicitly discloses all the limitations regarding claim 1, including the limitation of “substantially concurrently,” based on the interpretation of the claim language and the understanding (of the) prior art teachings, Appellants note that the broad interpretations made by the Examiner are **not consistent** with the specification and are **not consistent** with the interpretation of the specification that a person of ordinary skill in the art would make. As disclosed on page 2 (lines 16-26) of the original specification, “the present invention concurrently *segments an audio file and clusters the segments* corresponding to the same speaker.” Thus, the term “substantially concurrent” is related to the *parallel execution* of the segmentation and clustering steps. See, also, FIG. 2.

More specifically, Appellants note that these limitations are clearly captured in claim 1, which recites the limitations of identifying potential segment boundaries in said audio source; and clustering homogeneous segments from said audio

source substantially concurrently with said identifying step. Claim 1 requires the clustering of homogeneous segments *substantially concurrently* with said identifying step. Chen, therefore, actually teaches away from the present invention by teaching that the clustering is performed only after the audio stream has been segmented. Thus, contrary to the Examiner's assertion, the limitations cited by the Examiner in reference to Chen are *clearly relevant* to the consideration of claim 1.

Appellants also note that the references were not attacked individually, but were reviewed to demonstrate that *none* of the references contain the cited limitation required by the claims of the present invention and that, therefore, the prior art does not pose a bar to patentability.

Response to Final Office Action of December 18, 2007

The Examiner asserts that the claimed "clustering homogeneous segments...substantially concurrently with said identifying step (segmentation)" does not exclude the situation that "the clustering is performed only after the audio stream has been segmented."

Appellants note that the Examiner is including the case where the *clustering and segmenting are performed sequentially*. This is contrary to the claim requirement that the *clustering and segmenting are performed "substantially concurrently"*. The term "substantially concurrently" should be given patentable weight. In the cited example, however, the clustering and segmentation are performed sequentially; there is absolutely no degree to which the clustering and segmentation are performed "substantially concurrently." Thus, the Examiner's interpretation of the cited claims is *not a reasonable interpretation*.

Moreover, the *loop* illustrated in FIG. 2 demonstrates that the segmentation and clustering are performed substantially concurrently, as segmentation may be performed both *before and after clustering*.

The Examiner further asserts that "Chen's disclosure satisfies the claimed limitation under at least this minimum condition/assumption, because Chen recites 'comparing two models, one models the data as two Gaussian(s); the other models the data as just one Gaussian' to detect the changing point for segmentation (Chen, Section 3.1, page 4), such that at least two data groups (segments) are segmented (before) for

clustering speakers (Chen, Section 4, page 8).”

Appellants note that Chen describes a Gaussian distribution with means and variances. Chen assumes that the means of all the signals, $s_1 \dots s_k$, can be computed. This is only feasible if all the data required to compute the mean is available. For example, imagine there is a pipe conveying water from point A to point B. The observer at point B does not know how water will come over. (This is analogous to a radio signal or video stream.) In order to compute the mean volume of water emanating from a pipe per unit time, all the water can be collected into a container, the time required for the pipe to go dry can be measured, and the volume of water collected can be divided by the time required to collect it. This, in effect, is Chen's approach. In one aspect of the present invention, a running mean is used - that is, as and when the data arrives its statistics are computed. In terms of the cited example, the volume of water arriving every second (or some fixed multiple) is measured and running statistics are maintained. The result is therefore usable from the time the water starts emanating from the pipe.

In one embodiment of the present invention, the segments are automatically computed, where each segment is a speaker turn in a conversation. By gathering together all “similar” segments into a cluster, all of these speaker turns are recognized to correspond to one individual. This is done when the person finishes speaking his or her turn. In a roundtable of speakers, speech by the same speaker is able to be segmented and clustered as and when it occurs, versus after the entire roundtable has ended.

Chen, alternatively, outlines two clustering approaches in Sections 4 and 4.1. In the clustering approach of Section 4, the audio is first broken up into segments using the BIC criterion. The clustering begins after the entire audio has been broken into segments. In the real world, when dealing with real-time video or audio streams, the “entire” audio can be acquired, for example, only after one hits the ‘stop’ button on the recorder. After breaking the audio down into individual segments, Chen collects them into clusters. The number of clusters is open to begin with as is the cluster membership. Chen combines audio segments in different combinations in order to arrive at a globally optimized set of clusters as defined by the BIC criterion. As is stated in the first paragraph of Section 4.1, the process is very computationally expensive.

In the second “greedy” clustering approach (Section 4.1), Chen’s starting point is the same as the above. It is a set of individual audio segments realized “after” the entire audio stream has been broken up into segments. There is no simultaneous segmentation and clustering. Clustering follows segmentation; only the clustering algorithm is slightly optimized over the first technique. Chen states in the first sentence of section 4.1 that the process of clustering works by merging nearest nodes. Chen can do this only after he has access to “all” the nodes (each node corresponds to a segment). In line 3, paragraph 2 of Section 4.1, Chen teaches, “let $S = \{s_1, s_2, \dots, s_k\}$ be the current set of nodes ...” Here, Chen tacitly states that there is access to all the segments labeled s_1 through s_k , where k is the total number of segments (nodes), i.e., there is access to all the audio that is desired to be analyzed. In any real-time application, such as streaming audio or streaming video, there is access to all that has transpired thus far. Thus, segmentation and clustering can only be done based on past events.

Additional Cited References

Kleider et al. was also cited by the Examiner in rejecting claim 15 for its disclosure that the information of the speaker model data may include a speaker name. Appellants note that the inventors listed in United States Patent Number 5,157,763 (referred to by the Examiner in the Final Office Action) are not Kleider et al. Appellants did find, however, United States Patent Number 5,930,748 in the Notice of References Cited and respond to that reference below.

Appellants note that Kleider et al. is directed to a “method of identifying an individual from a predetermined set of individuals using a speech sample spoken by the individual. The speech sample comprises a plurality of spoken utterance, and each individual of the set has predetermined speaker model data.” Cited, Summary of the Invention. Kleider et al. do not address the issue of segmenting speech.

Thus, Kleider et al. do not disclose or suggest a “method of tracking a speaker in an audio source, said method comprising the steps of identifying potential segment boundaries in said audio source; and clustering homogeneous segments from said audio source substantially concurrently with said identifying step,” as required by independent claims 1, 16, 30, 31, 32 and 33 of the present invention. Similarly, independent claims 23, 34 and 35 require that the segmentation and clustering are

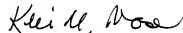
performed on the "same pass" through said audio source.

Conclusion

5 The rejections of the cited claims under sections 102 and 103 in view of
Chen, Kleider et al. or well known prior art, alone or in any combination, are therefore
believed to be improper and should be withdrawn. The remaining rejected dependent
claims are believed allowable for at least the reasons identified above with respect to the
independent claims.

10 The attention of the Examiner and the Appeal Board to this matter is
appreciated.

Respectfully,



15 Date: May 7, 2008

Kevin M. Mason
Attorney for Applicant(s)
Reg. No. 36,597
Ryan, Mason & Lewis, LLP
1300 Post Road, Suite 205
Fairfield, CT 06824
(203) 255-6560

20

APPENDIX

1. A method for tracking a speaker in an audio source, said method comprising the steps of:

- 5 identifying potential segment boundaries in said audio source; and
clustering homogeneous segments from said audio source substantially concurrently with said identifying step.

2. The method of claim 1, wherein said identifying step identifies segment
10 boundaries using a BIC model-selection criterion.

3. The method of claim 2, wherein a first model assumes there is no
boundary in a portion of said audio source and a second model assumes there is a
boundary in said portion of said audio source.

15 4. The method of claim 2, wherein a given sample, i , in said audio source is likely to be segment boundary if the following expression is negative:

$$\Delta BIC'_i = -\frac{n}{2} \log |\Sigma_w| + \frac{i}{2} \log |\Sigma_f| + \frac{n-i}{2} \log |\Sigma_s| \\ + \frac{1}{2} \lambda \left(d + \frac{d(d+1)}{2} \right) \log n$$

20

where $|\Sigma_w|$ is the determinant of the covariance of the window of all n samples, $|\Sigma_f|$ is the determinant of the covariance of the first subdivision of the window, and $|\Sigma_s|$ is the determinant of the covariance of the second subdivision of the window.

25

5. The method of claim 1, wherein said identifying step considers a smaller window size, n_s , of samples in areas where a segment boundary is unlikely to occur.

6. The method of claim 5, wherein said window size, n , is increased in a relatively slow manner when the window size is small and increases in a faster manner when the window size is larger.

5 7. The method of claim 5, wherein said window size, n , is initialized to a minimum value after a segment boundary is detected.

8. The method of claim 2, wherein said BIC model selection test is not performed at the border of each window of samples.

10

9. The method of claim 2, wherein said BIC model selection test is not performed when the window size, n , exceeds a predefined threshold.

10. The method of claim 1, wherein said clustering step is performed using a
15 BIC model-selection criterion.

11. The method of claim 10, wherein a first model assumes that two segments or clusters should be merged, and a second model assumes that said two segments or clusters should be maintained independently.

20

12. The method of claim 11, further comprising the step of merging said two clusters if a difference in BIC values for each of said models is positive.

13. The method of claim 1, wherein said clustering step is performed using K
25 previously identified clusters and M segments to be clustered.

14. The method of claim 1, further comprising the step of assigning a cluster identifier to each of said clusters.

30 15. The method of claim 1, further comprising the step of processing said audio source with a speaker identification engine to assign a speaker name to each of said

clusters.

16. A method for tracking a speaker in an audio source, said method comprising the steps of:

- 5 identifying potential segment boundaries in said audio source; and
clustering segments from said audio source corresponding to the same
speaker substantially concurrently with said identifying step.

17. The method of claim 16, wherein said identifying step identifies segment
10 boundaries using a BIC model-selection criterion.

18. The method of claim 17, wherein a first model assumes there is no
boundary in a portion of said audio source and a second model assumes there is a
boundary in said portion of said audio source.

15 19. The method of claim 16, wherein said identifying step considers a smaller
window size, n , of samples in areas where a segment boundary is unlikely to occur.

20. The method of claim 17, wherein said BIC model selection test is not
20 performed where the detection of a boundary is unlikely to occur.

21. The method of claim 16, wherein said clustering step is performed using a
BIC model-selection criterion, where a first model assumes that two segments or clusters
should be merged, and a second model assumes that said two segments or clusters should
25 be maintained independently.

22. The method of claim 16, wherein said clustering step is performed using K
previously identified clusters and M segments to be clustered.

30 23. A method for tracking a speaker in an audio source, said method
comprising the steps of:

identifying potential segment boundaries during a pass through said audio source; and

clustering segments from said audio source corresponding to the same speaker during said same pass through said audio source.

5

24. The method of claim 23, wherein said identifying step identifies segment boundaries using a BIC model-selection criterion.

25. The method of claim 24, wherein a first model assumes there is no boundary in a portion of said audio source and a second model assumes there is a boundary in said portion of said audio source.

10

26. The method of claim 23, wherein said identifying step considers a smaller window size, n , of samples in areas where a segment boundary is unlikely to occur.

15

27. The method of claim 24, wherein said BIC model selection test is not performed where the detection of a boundary is unlikely to occur.

28. The method of claim 23, wherein said clustering step is performed using a BIC model-selection criterion, where a first model assumes that two segments or clusters should be merged, and a second model assumes that said two segments or clusters should be maintained independently.

20

29. The method of claim 23, wherein said clustering step is performed using K previously identified clusters and M segments to be clustered.

25

30. A system for tracking a speaker in an audio source, comprising:
a memory that stores computer-readable code; and
a processor operatively coupled to said memory, said processor configured to implement said computer-readable code, said computer-readable code configured to:
identify potential segment boundaries in said audio source; and

30

cluster homogeneous segments from said audio source substantially concurrently with said identification of segment boundaries.

31. An article of manufacture, comprising:

5 a computer readable medium having computer readable code means embodied thereon, said computer readable program code means comprising:

a step to identify potential segment boundaries in said audio source; and

a step to cluster homogeneous segments from said audio source substantially concurrently with said identification of segment boundaries.

10

32. A system for tracking a speaker in an audio source, comprising:

a memory that stores computer-readable code; and

a processor operatively coupled to said memory, said processor configured to implement said computer-readable code, said computer-readable code configured to:

15

identify potential segment boundaries in said audio source; and

cluster segments from said audio source corresponding to the same speaker substantially concurrently with said identification of segment boundaries.

33. An article of manufacture, comprising:

20

a computer readable medium having computer readable code means embodied thereon, said computer readable program code means comprising:

a step to identify potential segment boundaries in said audio source; and

a step to cluster segments from said audio source corresponding to the same speaker substantially concurrently with said identification of segment boundaries.

25

34. A system for tracking a speaker in an audio source, comprising:

a memory that stores computer-readable code; and

a processor operatively coupled to said memory, said processor configured to implement said computer-readable code, said computer-readable code configured to:

30

identify potential segment boundaries during a pass through said audio source; and

cluster segments from said audio source corresponding to the same speaker during said same pass through said audio source.

35. An article of manufacture, comprising:

5 a computer readable medium having computer readable code means embodied thereon, said computer readable program code means comprising:

a step to identify potential segment boundaries during a pass through said audio source; and

10 a step to cluster segments from said audio source corresponding to the same speaker during said same pass through said audio source.

EVIDENCE APPENDIX

There is no evidence submitted pursuant to § 1.130, 1.131, or 1.132 or entered by the Examiner and relied upon by appellant.

RELATED PROCEEDINGS APPENDIX

There are no known decisions rendered by a court or the Board in any proceeding identified pursuant to paragraph (c)(1)(ii) of 37 CFR 41.37.